DOCUMENT RESUME

ED 462 412 TM 033 674

AUTHOR Zimmerman, Corinne; Glaser, Robert
TITLE Testing Positive versus Negative Claims: A Preliminary Investigation of the Role of Cover Story on the Assessment of Experimental Design Skills. CSE Technical Report.
INSTITUTION California Univ., Los Angeles. Center for the Study of Evaluation.; National Center for Research on Evaluation, Standards, and Student Testing, Los Angeles, CA.
SPONS AGENCY Office of Educational Research and Improvement (ED), Washington, DC.
REPORT NO CSE-TR-554
PUB DATE 2001-12-00
NOTE 19p.
CONTRACT R305B60002
AVAILABLE FROM UCLA/Center for the Study of Evaluation, 301 GSE&IS, box 951522, Los Angeles, CA 90095-1522 ($3.50). Tel: 310-206-1532. For full text: http://www.cse.ucla.edu/CRESST/Reports/TR554.pdf%20.
PUB TYPE Reports - Research (143)
EDRS PRICE MF01/PC01 Plus Postage.
DESCRIPTORS Ability; *Elementary School Students; *Experiments; Intermediate Grades; Research Design; *Test Items
IDENTIFIERS Open Ended Questions

ABSTRACT

This paper presents results from a preliminary investigation in which researchers manipulated the cover story of an open-ended assessment that required students to design an experiment. Participants were 27 sixth graders from a culturally diverse urban school. In one version of the cover story, students were asked to design an experiment to test a negative claim (i.e., that tap water is "bad" for plants). In a second version, they were to test a positive claim (i.e., that coffee grinds are "good" for plants). Differences between the two cover stories were found. In the negative condition, students were more likely to suggest controlled designs, to identify the correct independent variable, and to test the claim directly. In the positive condition, students proposed uncontrolled designs and did not select the correct variable to test. When asked to test a positive claim, students seemed to approach the task as if the goal was to test the generality of the claim (by selecting different plant types as the focal variable). Students of all ability levels were influenced by the cover story; however, the effect was greatest for students of low ability. Implications for assessment are discussed. Two appendixes contain the cover stories and the coding categories for design score for the task. (Contains 4 tables and 10 references.) (Author/SLD)

Reproductions supplied by EDRS are the best that can be made
from the original document.

CRESST
National Center for Research on Evaluation, Standards, and Student Testing

PERMISSION TO REPRODUCE AND DISSEMINATE THIS MATERIAL HAS BEEN GRANTED BY

K. Hurst

TO THE EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)

1

U.S. DEPARTMENT OF EDUCATION
Office of Educational Research and Improvement
EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)

☐ This document has been reproduced as received from the person or organization originating it.

☐ Minor changes have been made to improve reproduction quality.

• Points of view or opinions stated in this document do not necessarily represent official OERI position or policy.

ED 462 412

Testing Positive Versus Negative Claims: A Preliminary Investigation of the Role of Cover Story on the Assessment of Experimental Design Skills

CSE Technical Report 554

Corinne Zimmerman and Robert Glaser
CRESST/Learning Research and Development Center
University of Pittsburgh

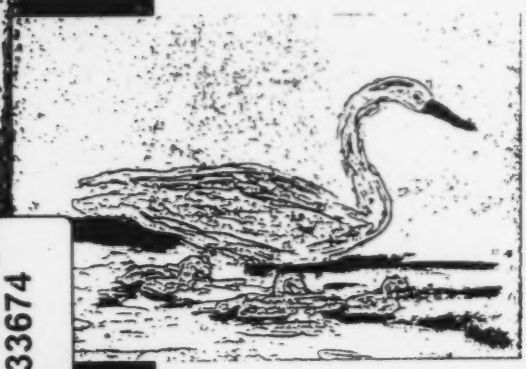

TM033674

UCLA Center for the Study of Evaluation

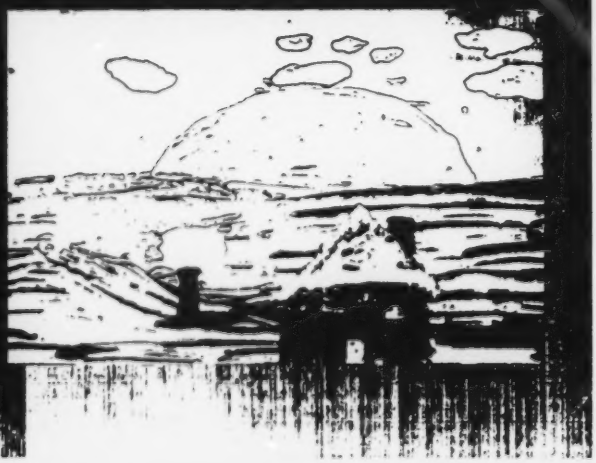

Testing Positive Versus Negative Claims:
A Preliminary Investigation of the Role of Cover Story
on the Assessment of Experimental Design Skills

CSE Technical Report 554

Corinne Zimmerman and Robert Glaser
CRESST/Learning Research and Development Center
University of Pittsburgh

December 2001

Center for the Study of Evaluation
National Center for Research on Evaluation,
Standards, and Student Testing
Graduate School of Education & Information Studies
University of California, Los Angeles
Los Angeles, CA 90095-1522
(310) 206-1532

Project 1.1. Models-Based Assessment Design: Individual and Group Problem Solving—Theory Into Practice. Robert Glaser, CRESST/Learning Research and Development Center, University of Pittsburgh, and Gail Baxter, Educational Testing Service, Project Directors.

Copyright © 2001 The Regents of the University of California

The work reported herein was supported under the Educational Research and Development Center Program, PR/Award Number R305B60002, as administered by the Office of Educational Research and Improvement, U. S. Department of Education.

The findings and opinions expressed in this report do not reflect the positions or policies of the National Institute on Student Achievement, Curriculum, and Assessment, the Office of Educational Research and Improvement, or the U. S. Department of Education.

TESTING POSITIVE VERSUS NEGATIVE CLAIMS: A PRELIMINARY INVESTIGATION OF THE ROLE OF COVER STORY ON THE ASSESSMENT OF EXPERIMENTAL DESIGN SKILLS [1]

Corinne Zimmerman and Robert Glaser
CRESST/Learning Research and Development Center
University of Pittsburgh

Abstract

There are many factors that potentially could influence student performance on assessments designed to tap reasoning and problem solving in science. We present results from a preliminary investigation in which we manipulated the cover story of an open-ended assessment that required students to design an experiment. In one version of the cover story, students were asked to design an experiment to test a negative claim (i.e., that tap water is "bad" for plants). In a second version, they were to test a positive claim (i.e., that coffee grinds are "good" for plants). Differences between the two cover stories were found. In the negative condition, students were more likely to suggest controlled designs, to identify the correct independent variable, and to test the claim directly. In the positive condition, students proposed uncontrolled designs and did not select the correct variable to test. When asked to test a positive claim, students seemed to approach the task as if the goal were to test the generality of the claim (by selecting different plant types as the focal variable). Students of all ability levels were influenced by cover story; however, the effect was greatest for students of low ability. Implications for assessment are discussed.

Developing assessments that measure problem solving and reasoning continues to be a challenging task in science education. Emphasis in classroom instruction is placed on both conceptual understanding and inquiry-based investigation skills because they reflect authentic scientific activity (e.g., American Association for the Advancement of Science, 1993; National Assessment Governing Board [NAGB], 1996; National Research Council, 1996, 2000). Assessments are needed to tap the concepts and reasoning skills that are elicited by such activities and that are consistent with these learning objectives.

---

[1] We acknowledge Laura Moin's assistance with the interviews and coding, and thank Nancy Lavigne and David Klahr for helpful comments on earlier drafts of the manuscript.

Baxter and Glaser (1998) have suggested that the effort to create open-ended assessment tasks can and should be informed by cognitive theory and research. Research on the characteristics of proficient problem solving can aid in the development of science assessments by relating these proficiencies to particular task characteristics. For example, the amount of content knowledge required to perform a task is a characteristic that may influence cognitive proficiencies such as problem representation or the ability to generate explanations. Likewise, tasks that are more or less open ended (e.g., student-generated procedures versus step-by-step instructions) will afford different opportunities to display cognitive proficiencies, such as strategy use and self-monitoring.

Task characteristics such as the amount of content knowledge or process skills have been shown to influence performance on science assessments (Baxter & Glaser, 1998). There are many factors that could potentially affect student performance when assessment tasks are designed to elicit reasoning and problem solving. Other cognitive research can be used to inform the development of assessments. In particular, a pervasive finding with laboratory tasks of reasoning and problem solving is the effect that a cover story has on performance. For example, problems often are easier to solve when they are presented within a concrete cover story than in an abstract one (for a review see Evans & Over, 1996).

Educational assessments may bear little resemblance to the laboratory tasks used previously to study reasoning. It can be argued, however, that factors shown to influence problem solving or reasoning *generally* should be considered when designing assessments to tap reasoning and problem solving in the science classroom. That is, even though laboratory tasks and assessment tasks are not precisely the same, there are similarities between the kinds of tasks used to study problem solving and those used to evaluate problem solving in assessment situations. These similarities warrant an investigation into whether, and to what extent, factors that influence performance on laboratory tasks may influence performance on assessment tasks. There are potentially important implications for the development of tasks by educators and assessment specialists if particular factors, such as cover stories, can be demonstrated to influence reasoning and problem solving.

A few studies suggest that examining the role of cover story on assessment tasks may be warranted. In a study relevant to problem solving in science (Tschirgi, 1980), adults and students in Grades 2, 4, and 6 were asked to select the variables

that needed to be manipulated to determine which variable was responsible for a problem outcome. Story problems were used in which two or three variables were involved in producing either a good or a bad outcome (e.g., baking a cake, making a paper airplane). In the cake-baking example there were three variables: type of sweetener (sugar or honey), type of shortening (butter or margarine), and type of flour (white or whole wheat). Participants were told that a story character believed that the honey was responsible for the (good or bad) outcome and asked how the character could prove this. Three options were provided: (a) bake another cake using the same sweetener (i.e., honey), but change the shortening and flour; (b) use a different sweetener (i.e., sugar), but the same shortening and flour; or (c) change all the ingredients.

The type of outcome (good or bad) influenced the strategy for selecting an experiment to produce evidence. When the outcome was "positive" (i.e., a good cake) participants in all age groups selected option *a* above, described as a "*hold* one thing at a time" strategy for manipulating variables. For "negative" outcomes (i.e., a bad cake) they selected option *b*, which is the more valid "*vary* one thing at a time" strategy. Tschirgi (1980) suggested that this result could reflect participants' experience with everyday problem-solving situations in which the goal is to reproduce positive effects and eliminate negative effects. If the goal or outcome is a factor that influences problem solving generally, then it could also influence tasks designed to assess problem solving in science.

In a study directly relevant to science, Schauble, Klopfer, and Raghavan (1991) demonstrated that students' perceived goal can affect performance on scientific reasoning tasks. Fifth- and sixth-grade students worked on the canal task (an investigation in hydrodynamics) and the spring task (an investigation of hydrostatics) under two different instructional cover stories. In each task, several factors contributed to the outcome. In the canal task, for example, students could experiment with canal depth, boat size, boat shape, and boat weight to determine which variables affect the travel time of boats. In one cover story, students were introduced to the task and asked to work as "scientists," and in the second cover story, they were asked to work as "engineers." When working as scientists, the instructional goal was to determine which factors made a difference and which ones did not. When working as engineers, the instructional goal was optimization, that is, to produce a desired effect (e.g., the fastest boat, in the canal task).

These two instructional cover stories resulted in different performances. In the scientist context, students worked more systematically, establishing the effect of each variable, alone and in combination. There was an effort to determine the factors that were causal as well as those that were not causal. In the engineering context, students selected highly contrastive combinations (e.g., a large, heavy boat versus a small, light boat) and focused on factors believed to be causal while overlooking factors believed or demonstrated to be noncausal. The approach to experimentation while students were acting as engineers was characterized as "try-and-see." When students were acting as scientists, the approach was characterized as "theory-driven." These findings support the idea that the perceived goal (i.e., understanding or optimization) can influence the manner in which students approach experimentation (i.e., try-and-see or theory driven).

In summary, task characteristics such as cover story have been shown to influence students' scientific reasoning on both paper-and-pencil tasks and open-ended experimentation tasks. These studies highlight the importance of considering the effects of cover story or task variations when designing instructional and assessment tasks that foster and elicit students' reasoning when they are engaged in scientific investigation.

In this paper, we focus on the influence of cover story on sixth-grade students' performance on an open-ended task administered upon completion of the *Experiments with Plants* curriculum unit (National Science Resource Center [NSRC], 1992). The goal of this curriculum unit is "to teach students how to design and conduct controlled investigative experiments" (p. 1). Two objectives guided the design of the present study. Our first objective was to determine whether or not students *trained* in process skills (e.g., identifying variables, planning experiments, and interpreting results) are susceptible to cover story effects. In previous research (e.g., Schauble et al., 1991; Tschirgi, 1980) students were not specifically trained in experimental design. We were interested to see whether performance differences for two different cover stories would occur in a group of students who spent 8 weeks learning about experimentation. Failure to find a cover story effect would indicate that even minimal training is sufficient to overcome particular "biases" reported in the literature, whereas a difference would suggest that the effect is pervasive despite instruction in experimental design.

Our second objective was to find out whether cover stories influenced performance for different groups of students. In particular, we were interested in

determining whether the effect, if found, would hold for students of differing ability levels in science (high, middle, low). One possibility, for example, was that high-ability students would not be influenced by superficial variations in cover story.

Method

Twenty-seven sixth-grade students from a culturally diverse school in an urban center participated in the study (9 boys, 18 girls). The teacher was asked to provide ratings of the students according to ability level in science. There were 11 high-ability, 10 middle-ability, and 6 low-ability students. Students were interviewed individually and instructed to "think aloud" while they performed a task that required them to design an experiment with plants. This task was intended to represent a typical end-of-unit assessment for the *Experiments with Plants* curriculum unit and was based on a task used previously in this school district (Raghavan, 1999). Students were instructed to respond to the following: (a) Describe and explain how you would set up the experiment; (b) describe what you would measure; and (c) design a table to record the data collected throughout your experiment.

Two different versions of the plants task were created (see Appendix A). In the *positive claim* version, students were asked to design an experiment to find out whether coffee grinds are "good" for green bean plants. In the *negative claim* version, students were asked to design an experiment to find out whether tap water is "bad" for green bean plants. Students were assigned randomly to one of these conditions. The negative claim group (n = 14) consisted of 5 high-, 6 middle-, and 3 low-ability students. The positive claim group (n = 13) consisted of 6 high-, 4 middle-, and 3 low-ability students. The two groups, therefore, had reasonably equivalent distributions of student ability levels.

A scoring scheme was devised to examine several aspects of students' overall performance on the task (see Appendix B). Credit was given whenever information within the following categories was mentioned: (a) questions and/or hypotheses (maximum 4 points), (b) materials, such as pots, soil, etc. (maximum 7 points), (c) experimental design features, including such things as background research, data recording, measures, etc. (maximum 21 points), and (d) data table design (maximum 8 points). A *design score* was calculated by tallying the number of different features from these four subcategories (maximum of 40 points). This measure is a frequency count of the range of research features that could be mentioned when designing an experiment. That is, there is no weighting of each item to reflect that, for example,

mentioning "record data daily" could be considered more important in the context of experimental design than "get soil and pots."

A second scoring scheme was used to assess the specific skills considered to be important when learning about the scientific process (e.g., NSRC, 1992). A *process score* was calculated that included the following five components: (a) manipulating only one variable, (b) manipulating the correct variable, (c) keeping conditions constant (i.e., controls such as the amount of sunlight, soil, water, or coffee grinds), (d) use of repeated measurements (e.g., use of trials, averages, or multiple plants per condition), and (e) systematic observation (i.e., measurements made on a regular basis over time). Students received 0-2 points for each process skill and an overall process score (maximum of 10 points). This measure has been used in previous research to evaluate the science curriculum in this school district (Raghavan, 1999).

Each protocol was coded by one individual. A second individual coded 25% of the sample (7 protocols). The initial agreement for the components of the design score was 95.5%. All "disagreements" were in the form of omissions. For the process score, initial agreement was 91.4% and disagreements were resolved through discussion.

Results

A 2 (cover story) x 3 (ability level) analysis of variance was used to analyze the design score and the process score. There was a significant difference between the two cover stories on the design score, with the negative claim group mentioning 18.64 design features on average, and the positive claim group mentioning 14.15 features, $F(1,21) = 4.91$, $p = .038$. The three ability groups did not differ on this measure, and there was no interaction between ability level and cover story.

For the process score, there was a main effect of cover story, $F(1, 21) = 13.01$, $p = .002$. The negative claim group had an average of 7.36 on the total process score (maximum of 10 points), and the positive claim group averaged 4.77. There was a main effect of ability level, $F(2, 21) = 3.43$, $p = .05$, with the low-, middle-, and high-ability groups scoring 4.3, 7.3, and 6.0, respectively. Post-hoc comparisons indicate a significant difference between the middle-ability group and the low-ability group, $F(1,21) = 6.85$, $p = .016$, but not between the high- and low-ability groups. These main effects must be considered in the context of the significant interaction between cover story and ability level, $F(2, 21) = 5.18$, $p = .015$. As can be seen in Figure 1, this

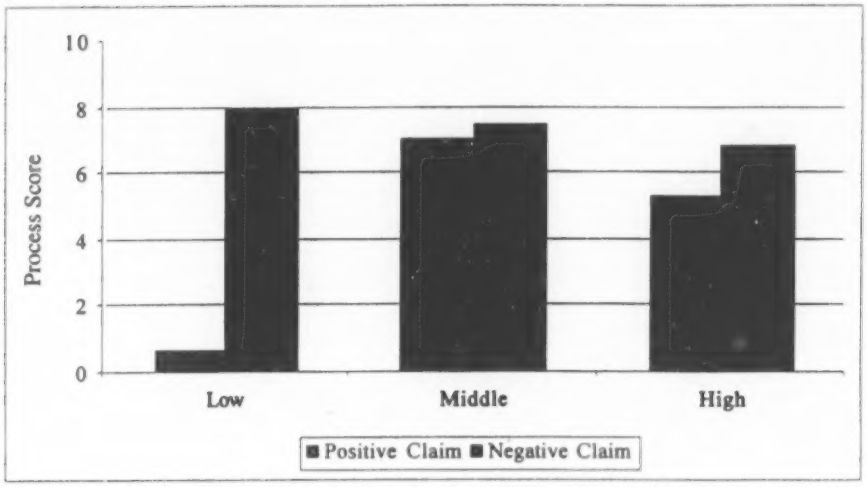

Figure 1. Mean process score by ability level and cover story.

interaction is due to the large difference between the positive claim and negative claim conditions for low-ability students, $F(1,21) = 17.6$, $p < .0001$.

Recall that the process score is comprised of five different skills. The percentage of students receiving credit for each of these five skills appears in Table 1. Statistical analyses for several of the individual process skills are not possible due to low expected frequencies. An analysis of the individual process skills by ability levels was not possible for the same reason.

Table 1

Percentage of Students in Each Condition Receiving Credit for Five Process Skills

Process skill	Positive claim	Negative claim
Manipulate the correct variable	23	79
Manipulate only one variable	23	100
Use of repeated measures	54	86
Systematic observation	54	71
Keep conditions constant (controls)	15	21

Each of the students' designs was characterized based on the main variable that was manipulated. The designs used by the 14 students in the negative claim condition are shown in Table 2. The majority of students selected either a simple design contrasting tap water with rain water, or a design comparing tap water, rain water, and other kinds of waters (e.g., sugar water, bleach water). A small number of students selected "plant type" as the manipulated variable, with only one of three explicitly mentioning the idea of using control plants.

The designs proposed by the 13 students in the positive claim condition are presented in Table 3. Only three students focused on coffee as the manipulated variable (recall that the goal was to test the claim that coffee grinds are "good" for plants). More than one half of the students focused on the type of plant as the manipulated variable (e.g., azaleas, green beans, roses, etc.). One fifth of the students in this group suggested designs that were not clear or were off-task.

Table 2

Number of Students Using Different Experimental Designs in the Negative Claim Condition

Experimental design	Number of students
Tap versus rain water	7
Type of water (tap, rain, others)	4
Type of plant plus control plant(s)	1
Type of plant (no control plants)	2

Table 3

Number of Students Using Different Experimental Designs in the Positive Claim Condition

Experimental design	Number of students
Coffee versus no coffee	1
Types of coffee plus no coffee	2
Type of plant plus control plant(s)	1
Type of plant (no control plants)	6
No obvious manipulated variable/design not clear	2
Other design[a]	1

[a] Student suggested putting green been seeds on plants to see effect.

Students' designs were further classified as either controlled or uncontrolled for both conditions regardless of manipulated variable. More students in the negative claim condition suggested controlled designs (12/14) than in the positive claim condition (4/13), and this relationship was statistically significant ($\chi^2 = 8.43$, $df = 1$, $p = 0.004$). The numbers of students with controlled and uncontrolled designs by cover story and ability level are presented in Table 4, but statistical tests are not reliable because of low expected frequencies. The most evident pattern emerging from this study is that across all ability levels, students in the positive group tended to suggest uncontrolled designs whereas those in the negative group suggested controlled designs.

Discussion

Our first objective was to determine whether cover story effects would occur in a group of students trained in a curriculum unit in which one of the key goals was to learn how to design and conduct experiments. The valence of the cover story affected aspects of students' performance on the end-of-unit assessment requiring the design of an experiment. The negative claim group scored higher on the design score and its various categories (i.e., questions/hypotheses, materials, experimental design features, data table). That is, students in the negative claim group mentioned more features of research overall than students in the positive claim group.

The most striking difference between the cover story groups was the superior ability of the negative claim group to manipulate the correct variable and to focus on manipulating only one variable at a time. In the negative claim group, students were instructed to design an experiment to test the claim that "tap water is bad for plants," and they tended to identify water as the focal variable. Students testing the

Table 4

Numbers of Students Proposing Controlled and Uncontrolled Designs by Ability Level and Cover Story (Positive vs. Negative)

	Positive claim		Negative claim	
Ability level	Controlled	Uncontrolled	Controlled	Uncontrolled
High	2	4	4	1
Middle	2	2	5	1
Low	0	3	3	0

positive claim that "coffee is good for plants" tended to identify plant type as the focal variable.

Our second objective was to find out whether cover story would differentially affect students of different ability levels. There was an overall effect of cover story, but it needs to be considered in the context of the interaction with ability level. Although students of all abilities performed better in the negative claim condition, there was a large difference between the negative and positive claim groups for students judged by their teacher to be of low ability. For the five process skills of interest, the low-ability students appear to be more susceptible to the effect of cover story.

A key learning objective for the curriculum unit was the ability to design controlled experiments. Students of all ability levels in the positive claim group were more likely to propose uncontrolled designs, whereas students of all ability levels in the negative claim group were more likely to propose controlled designs. This finding provides some evidence for the assertion that a simple variant in cover story has an impact on all students' performance.

In this study, the cover story within which a problem was presented influenced students' experimental designs. There are a number of potential explanations for why the cover story effect occurred, including (a) the valence of the claim, (b) differences in students' familiarity with the variables in the two cover stories, and (c) cueing differences in the cover stories.

The two cover stories differed in the valence of the claim in that the variable of interest was either "good" or "bad" for plants, and students were instructed to design an experiment to test this assertion. Tschirgi (1980) demonstrated that this characteristic of word problems influenced children's and adults' evaluation of experimental designs and suggested that individuals selected their strategies to reproduce positive effects and eliminate negative effects. The present results seem consistent with her findings and interpretation. Students in the two groups seemed to use different strategies for designing their experiments. Instead of testing the claim that coffee grinds are good for plants by setting up an experiment with green bean seeds (as suggested in both cover stories), students in the positive claim group acted as though they were *accepting* the claim and that their goal was to test the *generality* of it by using coffee grounds with a variety of plants (e.g., green beans, roses, pine trees, etc.). Proposing designs with plant type as the main independent

variable is consistent with the idea that they were attempting to reproduce the "good effect" of coffee grinds with different types of plants.

If there is a general tendency based on real-world problem solving to reproduce positive effects and eliminate negative effects, then any assessment that involves problem solving couched in a non-neutral way may either underestimate or overestimate students' performance. Based on the current findings, cover stories that have a negative valence may overestimate students' ability to design experiments whereas cover stories with a positive valence may underestimate students' abilities. There may be a tendency to accept positive claims, which could then lead students to misinterpret the goal of the assessment exercise. Schauble et al. (1991) demonstrated that students' perceived goal influenced how they approached open-ended experimentation tasks.

A second explanation for the performance differences has to do with the familiarity of the focal variables in the two cover stories. Water may be a more *familiar* variable than coffee grinds. There are a number of possible variables that students can use in their class projects with plants, but the teacher's manual specifies that teachers should dissuade students from manipulating water (NSRC, 1992). Nonetheless, students' understanding of the role of water may make this an inherently more familiar variable than coffee grinds in the plant context, and therefore needs to be considered as a potential explanation for the performance differences found in the present study.

A third possible explanation concerns the fact that the two cover stories were not as isomorphic as they could have been. Our intention was to make them as similar as possible, while at the same time being plausible. For example, it seemed implausible to create a negative claim version in which Sonia's grandmother (one of the story characters) asserts that coffee grinds are "bad" for plants (i.e., if she thought they were bad, realistically, she would not be using them on the azaleas). The claim that tap water was "bad" seemed plausible, and most students, when asked to predict the outcome in this condition, thought that it was a reasonable claim because tap water contains lead and other chemicals.[2] As can be seen in Appendix A, in order to make the claim realistic, Sonia's grandmother is using rainwater instead of tap water. The mention of an additional type of water may have

---

[2] Analyses of the prediction data were not presented because a second interviewer neglected to ask this question on the first day of testing.

been a cue for students to focus on water and not plant type. This same type of cue is not present in the cover story with coffee grinds and must be considered as a possible explanation for why the students tended to focus on plant type as the main variable instead of coffee grinds.

Despite the fact that there are several possible explanations for the performance differences, the fact that a slight variation in cover story had an impact has implications for both assessment and instruction. Desired characteristics and attributes of assessment exercises are outlined in the NAEP Science Consensus project (NAGB, 1996), including the idea that "problems need to be placed in new contexts, applied to new situations, or have new elements introduced that preclude students from simply recalling what they have done before" (p. 33). None of the students in the present study conducted class projects with either water or coffee grinds, but both cover stories represent likely problem scenarios to which students might be expected to generalize their knowledge.

General Discussion

One goal of the present study was to determine whether cover stories with positive versus negative claims would result in performance differences on an end-of-unit assessment. Differences in particular process skills were found between the negative claim version (i.e., whether tap water is "bad" for green bean plants) and the positive claim version (i.e., whether coffee grinds are "good" for green bean plants). Ideally, students should be able to transfer the inquiry skills learned in the classroom regardless of cover story. That is, even a "near transfer" task proved difficult for some students depending on the nature of the cover story used. A student's predisposition to be influenced by superficial task characteristics needs to be explicitly addressed in the classroom and in assessment design. Dealing with this issue is critical given that the positive claim version of the design task is currently used to assess students' performance at the end of the *Experiments with Plants* curriculum unit (Raghavan, 1999). These types of assessment (i.e., with cover story variations) would better serve the needs of students if they were included as part of the curriculum, either as pedagogical examples or as embedded assessments that could be used to inform and modify subsequent instruction.

An implication of this study is that students require a broad range of instructional experiences to generalize their developing skills. Diverse activities are needed to obtain information about students' knowledge and to provide learners

with opportunities to develop flexibility in their thinking. This effect needs to be explicitly addressed in the classroom through discussions that focus on the issues of experimental outcomes and biases. Multiple examples of each form (i.e., positive, negative) can be provided as a basis of discussion and further instruction. Identifying students' existing knowledge and building on this knowledge through activities appears to be the key to facilitating the development of scientific investigation skills.

References

American Association for the Advancement of Science. (1993). *Benchmarks for science literacy*. New York: Oxford University Press.

Baxter, G. P., & Glaser, R. (1998). Investigating the cognitive complexity of science assessments. *Educational Measurement: Issues and Practice, 17*, 37-45.

Evans, J. St. B. T., & Over, D. E. (1996). *Rationality and reasoning*. Hove, UK: Psychology Press.

National Assessment Governing Board. (1996). *Science framework for the 1996 National Assessment of Educational Progress*. Washington, DC: Author.

National Research Council. (1996). *National science education standards*. Washington, DC: National Academy Press.

National Research Council. (2000). *Inquiry and the national science standards*. Washington, DC: National Academy Press.

National Science Resource Center. (1992). *Experiments with plants: Teacher's guide*. Burlington, NC: Carolina Biological Supply.

Raghavan, K. V. (1999). *Local systemic change: 1998-99 evaluation report for ASSET Inc.* Pittsburgh, PA: University of Pittsburgh, Learning Research and Development Center.

Schauble, L., Klopfer, L. E., & Raghavan, K. V. (1991). Students' transition from an engineering model to a science model of experimentation. *Journal of Research in Science Teaching, 28*, 859-882.

Tschirgi, J. E. (1980). Sensible reasoning: A hypothesis about hypotheses. *Child Development, 51*, 1-10.

Appendix A

Cover Stories for the Plants Task

Positive Claim Cover Story

Sonia, a fourth grader, got curious as she watched her grandmother empty the coffee grinds from her coffee machine onto the base of an azalea bush. When asked why, her grandmother said that used coffee grinds are "good" for the azaleas. This gave Sonia an idea. She wanted to do some experiments to find out if coffee grinds are "good" for other plants. She collected some green bean seeds for her experiments. Sonia asked you to help her set up the experiments.

Negative Claim Cover Story

Sonia, a fourth grader, got curious as she watched her grandmother water an azalea bush with water she got from a rain barrel. When Sonia asked her grandmother why she didn't just use the hose to water the bushes, her grandmother said that tap water is "bad" for the azaleas. This gave Sonia an idea. She wanted to do some experiments to find out if tap water is "bad" for other plants. She collected some green bean seeds for her experiments. Sonia asked you to help her set up the experiments.

Appendix B

Coding Categories for Design Score for the Plants Task

Scoring is 0/1 for each unless otherwise noted.

1. Hypothesis or Question (4 possible points)
   - Mentions "hypothesis"
   - Specific mention of what the hypothesis is
   - Mentions "question"
   - Specific mention of what the question is

2. Materials (7 possible points)
   - Generic term "materials"
   - Pots
   - Seeds
   - Soil
   - Water
   - Sun/light
   - Other: specify (e.g., a "helper")

3. Experimental Design Features (21 possible points)
   - Use of constants (e.g., same amount of coffee/water, location, light, start/finish time, etc.)
   - Controlled design (0-2)
   - Multiple trials
   - Generic term "procedures"
   - Generic term "plan/s"
   - Background research (library, Internet)
   - *Start* the experiment (differentiation between planning and starting emphasized in curriculum)
   - Data recording
   - Record data *every day* or frequently (i.e., > once)
   - Notebook
   - Graphing
   - *Number* of measures (acceptable based on course materials: height, pods, leaves, flowers, color, health, life span, cotyledons, rate of growth, germination day, first days, etc.)
   - Take averages

4. Data Table Design (8 possible points)
   - Title
   - Labels
   - Multiple trials
   - Average
   - Date
   - Table sections for experimental versus control
   - Table consistent with experimental design (0-2)